BEFORE THE
Federal Communications Commission
WASHINGTON, D.C.

In the Matter of	
Preserving the Open Internet	GN Docket No. 09-191
Broadband Industry Practices	WC Docket No. 07-52

COMMENTS OF BRIGHT HOUSE NETWORKS, LLC
IN RESPONSE TO SUPPLEMENTAL NOTICE

Marva Johnson
Vice President
Technology Policy and Industry Affairs
Bright House Networks, LLC

Paul Glist
Christopher W. Savage
Davis Wright Tremaine LLP
1919 Pennsylvania Ave. NW
Washington, D.C. 20006-3402

Arthur J. Steinhauer
Cody Harrison
Sabin, Bermant & Gould LLP
Four Times Square
New York, NY 10036

Counsel for Bright House Networks, LLC

October 12, 2010

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The Supplemental Notice[1] seeks additional comments and information on the role of "specialized services" in potential regulatory approaches to maintaining an open Internet, and on the issue of extending open Internet principles to mobile wireless platforms.

The Commission should not impose any regulatory obligations with respect to specialized services at this time. End users have been receiving rapidly increasing Internet access capacity via multifunction last mile facilities operated by cable operators, telephone companies, and Commercial Mobile Radio Service ("CMRS") licensees. From this perspective, there is no evidence that there is any specialized services "problem" that needs to be solved. More broadly, the networking community – including retail broadband networks, backbones, private content delivery optimization networks, data center providers, etc. – has a long history of working cooperatively to develop suitable ways to deliver content and applications with particular technical or business requirements. For example, with little fanfare, various content delivery networks have been deployed to take large amounts of traffic off the Internet backbone and deliver it directly, often on a settlement-free basis, to retail networks, in order to avoid problems with high latency. This cooperative approach arises because the value of the network to end users is produced by the system as a whole, not by any one member. As a result, all members have an interest in managing the jointly-created resource in an efficient manner. In addition, the technical complexity of multiple entities coordinating to provide an end-to-end service is so great that the only realistic way to approach these problems is in an informal cooperative manner. A variety mechanisms are already in place to allow this process to occur. It can and will address any issues regarding specialized services that may arise.

---

1 "Further Inquiry into Two Under-Developed Issues in the Open Internet Proceeding," DA 10-1667, GN Docket No. 10-191, WC Docket No. 07-52 (September 1, 2010) ("Supplemental Notice").

With respect to wireless, the Commission should not treat a network operated by a CMRS licensee differently from any other retail broadband Internet access network with regard to open Internet principles or rules regarding specialized services. The distinction between "wireless" and "wired" networks has become entirely artificial. Wireless devices – from laptop and notebook computers, to mobile "smart phones," to iPads – are the normal means by which consumers obtain broadband Internet access. These devices all initially communicate with a nearby antenna, but then immediately transfer their Internet communications to a "wired" network. Moreover, CMRS-based devices like the iPhone and many Android-based phones, are, increasingly, WiFi devices as well, allowing users to switch between CMRS and WiFi networks at will, further blurring any distinction that may once have existed. In addition, every retail broadband network, regardless of technology, must use a number of strategies to manage network capacity to meet growing end user demands. Furthermore, as end users increasingly rely on wireless devices for Internet access, excluding wireless networks from the Commission's policies regarding open Internet access – including specialized services – would senselessly leave substantial public interest benefits on the table. As a result, the same rules and polices should apply to all retail broadband Internet access networks.

BEFORE THE
Federal Communications Commission
WASHINGTON, D.C.

In the Matter of	
Preserving the Open Internet	GN Docket No. 09-191
Broadband Industry Practices	WC Docket No. 07-52

COMMENTS OF BRIGHT HOUSE NETWORKS, LLC
IN RESPONSE TO SUPPLEMENTAL NOTICE

Bright House Networks, LLC ("Bright House Networks") submits these comments in response to the Supplemental Public Notice released in this matter on September 1, 2010.[2]

I. MULTI-FUNCTION NETWORKS ARE THE RULE, NOT THE EXCEPTION.

Before addressing the specific questions posed in the Supplemental Notice, it is important to address what may be a misconception regarding the physical networks used to deliver broadband Internet access services.

The Supplemental Notice states that the Commission is aware that "broadband providers ***may*** provide other services over the same last-mile facilities used to provide broadband Internet access service," and then identifies certain potential issues that might arise were that condition to occur.[3] This statement of the issue is misleading: the overwhelming majority of broadband Internet access is provided via network facilities that not only provide other services, but that have, from the outset, been designed to do so. Multifunction networks are – and always have been – the rule, not the exception.

---

2 "Further Inquiry into Two Under-Developed Issues in the Open Internet Proceeding," DA 10-1667, GN Docket No. 10-191, WC Docket No. 07-52 (September 1, 2010) ("Supplemental Notice").

3 *Id.* at 2 (emphasis added).

The overwhelming dominance of multifunction retail networks is one result of the "convergence" of different legacy networks that has been noted for decades.[4] Part of what convergence ***means*** is that the same physical network facilities – the "last-mile facilities" singled out in the Supplemental Notice – that might once have provided one or a few services now offer many.[5] While it may still be common to label a network based on the legacy service that it was originally designed to provide, today essentially every retail broadband network provides many services over the same last-mile facilities.[6]

For this reason, the Commission should not assume that the default case is a network dedicated to and optimized for the delivery of broadband Internet access, with other services attached later as appendages. And the Commission should not assume that the default case is that all services – voice, video, and data – will be delivered via broadband Internet access. All of the relevant physical networks were built and designed, at least initially, for things ***other than*** broadband Internet access. And, as discussed below, there are any number of reasons why it

[4] As early as 1991, the Commission began to recognize the "increasing convergence of previously separate markets embracing voice, data, graphics and video." *See* In the Matter of Telephone Company-Cable Television Cross-Ownership Rules, *Further Notice of Proposed Rulemaking, First Report and Order and Second Further Notice of Inquiry*, 7 FCC Rcd 300 (1991) at ¶ 7. The Commission now recognizes that we have entered "the modern era of converged digital services." In the Matter of Framework for Broadband Internet Service, *Notice of Inquiry*, 25 FCC Rcd 7866 (2010) at ¶ 3.

[5] Supplemental Notice at 2 (stating that "the relationship between open Internet protections and services that are provided over the same last-mile facilities" presents a "complex issue"); *id.* (noting that providers may provide "other services over the same last-mile facilities used to provide broadband Internet access service").

[6] We recognize that some network providers – for example, Clearwire – are deploying networks that are essentially dedicated to broadband Internet access. Such networks, however, are the exception that proves the general rule that Internet access is overwhelmingly provided over multi-function networks. Moreover, different retail network operators may take very different technical approaches to offering multiple services on the same physical last-mile facilities. For example, an operator may use logically separate "networks" to deliver different services. *See, e.g.,* J. Finn, "PON Technology in the Verizon Network," Proceedings of IEEE GLOBECOM 2008, *available online at:* http://202.194.20.8/proc/GLOBECOM2008/DATA/13-07-04.PDF at page 2 (describing Verizon use of separate wavelengths for video transmissions versus voice and data transmissions). Another operator might choose to integrated multiple services into a single bitstream. Obviously, the Commission should not base regulatory distinctions or policy on these kinds of highly technical network engineering choices.

makes sense to offer certain new and existing service using something other than the best-efforts Internet.

With this perspective, we turn to the specific questions raised in the Supplemental Notice.

II. THE DEPLOYMENT OF "SPECIALIZED SERVICES" WILL NOT SUPPLANT OR RESTRICT BROADBAND INTERNET ACCESS.

The Supplemental Notice asks whether network operators "may constrict or fail to continue expanding the network capacity allocated to broadband Internet access service in order to provide more capacity for specialized services."[7] This is the key concern in the Supplemental Notice. If the network capacity available for broadband Internet access is large and growing – which it is – any other concerns are speculative. The need to meet continuously growing end user demand for additional broadband capacity, as well as competition from other retail network operators, has ensured that there is no shortage of capacity devoted to broadband Internet access service. As a result, Commission action with respect to specialized services is not needed.

A. Bright House Networks Has Consistently Invested To Increase The Network Capacity Devoted To Broadband Internet Access.

As noted above, network operators did not design and build their networks to provide broadband Internet access, and then add on other services – a formulation that evokes images of network capacity being clawed away from Internet access service to meet the needs of non-Internet services.[8] Any such image is profoundly misleading. Cable operators, for example, have networks that were initially designed to provide cable service. Then, as the demand for broadband Internet access became evident, cable operators began modifying and expanding their networks, with more and more capacity devoted to broadband Internet access. Broadband Internet access service is ***growing*** on our network – not shrinking. For our part, Bright House

---

[7] Supplemental Notice at 2.

[8] *See, e.g.,* Center for Democracy & Technology Comments at 49 (referring to a supposed "risk … of gradual erosion" of broadband Internet capacity).

Networks is making investments in new network capacity measured in the hundreds of millions of dollars to ensure that that growth can continue.

The main driver for the growth in Internet usage has been the inventiveness of content and application providers, who develop interesting and useful websites and applications that require higher throughput than was previously available. Ten years ago, the average web page could be easily viewed with a connection speed of about 128 Kbps, while today it takes almost 30 times the bandwidth – a connection speed of about 3.5 Mbps – to download the average web page rapidly enough to avoid end user dissatisfaction with delays:[9]

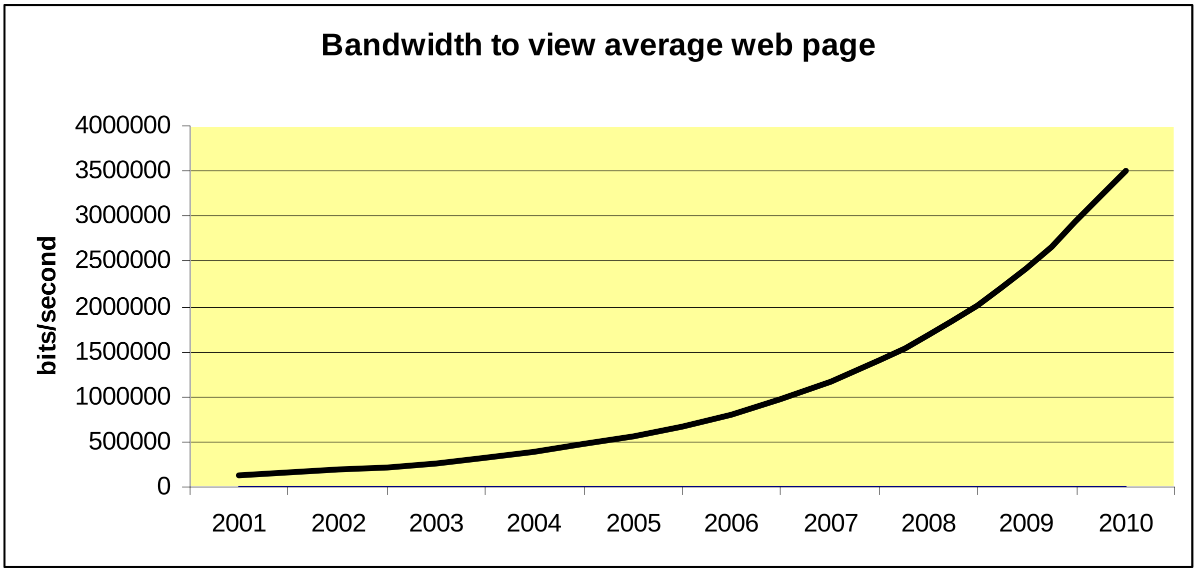

---

[9] The information in the following chart, and the others in these comments, is based on information supplied by Bright House Networks' engineering staff, who are familiar not only with Bright House Networks' specific operations, but also overall industry trends. We have interpolated smooth growth between the end points rather than tried to generate specific year-by-year data. With respect to the average size of a web page underlying the chart below, in 2003 the average web page contained only about 93.7 kilobytes of data. *See* "Average Web Page Size Quintuples Since 2003," http://www.websiteoptimization.com/speed/tweak/average-web-page/. At this size, an average page could be transmitted in about 5 seconds at a data rate of 128 Kbps. (Note that 100 kilo***bytes*** of data translates to 800 kilo***bits*** to be transmitted, because each byte contains 8 bits.) By 2009 that had grown to more than 500 kilobytes, or 4 megabits, of data (*see id.*) – with consumers, all the while, becoming less tolerant of "slow" downloads and delays in receiving the information they want. Other factors affecting download speeds include the number of external objects contained in the average web page that need to be separately retrieved, which grew by 2.5 times between 2003 and 2009. *Id.*

This increase in the bandwidth needed to promptly view the average web page is one aspect of the changing nature of the applications end users are accessing. Where once “the Internet” meant nothing more than text-based email and downloading static, text-based web pages, now “the Internet” means a wide array of complex, graphics- and video-intensive interactive functionalities. Moreover, ten years ago, today’s common “streaming” applications – for example, music services such as Pandora, and video services such as YouTube – did not exist, or were only in their infancy. Today, use of these services is ubiquitous and growing. Online gaming, which once required symmetrical 64 Kbps bandwidth to work – roughly the capacity of a digital voice line – now requires at least 256 Kbps in both directions. All of these factors have created an end user expectation of ever-increasing bandwidth.[10]

Keeping up with end user demands has required enormous and varied investments by Bright House Networks and other network operators. For example, residential end users can now download data at rates of 100 Mbps (and in some cases higher), an improvement by a factor of 100 from a more common maximum effective rate of 1 Mbps ten years ago:

[10] Broadband Internet access is an “experience” good. That is, the more that end users have access to and use broadband Internet access – and see what it can do – the more they value it and the more of it they want. *See* G. L. Rosston, S.J. Savage & D. M. Waldman, *Household Demand for Broadband Internet In 2010,* 10 Berkeley Electronic Journal of Economic Analysis & Policy, Issue 1, Article 79 (2010). Ultimately, though, what makes this occur – that is, the ***reason*** broadband Internet access is an “experience” good – is the interesting, useful, compelling, and clever content and applications end users can find using the Internet.

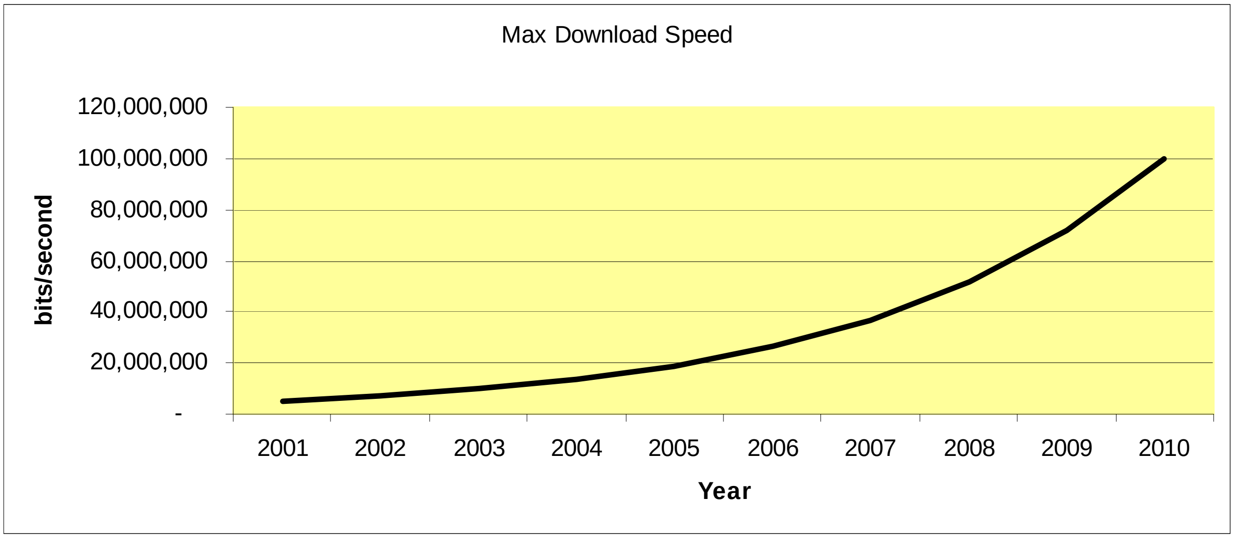

The increase in maximum download capacity, combined with the vastly greater data requirements of the average web page, has resulted in very large increases in average monthly data usage. Indeed, the average monthly usage of cable broadband Internet subscribers has increased a thousand-fold, from roughly 10 megabytes of data per month ten years ago, to approximately 10 ***gigabytes*** per customer per month today:

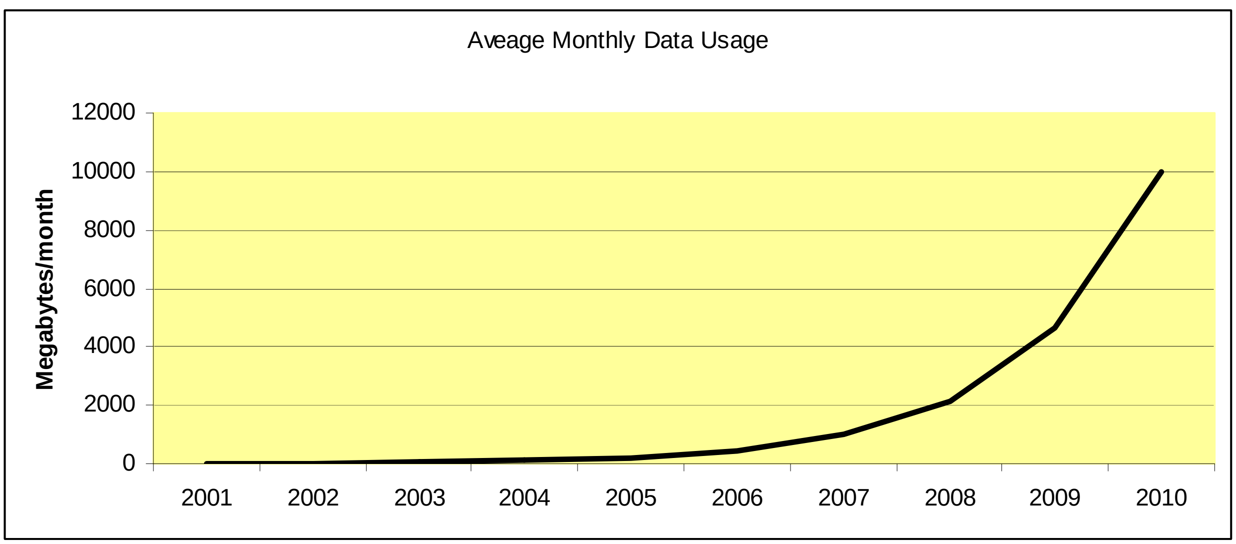

These increases in maximum download capacity and average monthly data usage have been made possible through continuing investments in network capacity and functionality. Focusing on Bright House Networks:

- We have deployed thousands of strand miles of optical fiber from our head-ends out to neighborhood nodes in order to enable the delivery of greater bandwidth to each end user over the coaxial cable portion of our plant.

- Our system has continuously upgraded from pre-DOCSIS arrangements, to DOCSIS 1.0, to DOCSIS 2.0 to, now, DOCSIS 3.0 in order to increase the functionality of the overall transmission and delivery of data. For example, with DOCSIS 3.0 we can bond multiple channels for improved capacity. DOCSIS 3.0 also enables greater reliability, with fewer dropped packets and other technical improvements, such as support for IPv6 Internet addressing.[11]

- We have improved the modulation techniques on our systems in order to increase our hybrid fiber-coax capacity and increase the number of system channels devoted to data services. This has increased the downstream capacity from 25 Mbps to 160 Mbps in the last 10 years. 320 Mbps is projected to be available in the next 3 years.

- With regard to upstream transmissions, our efforts have increased the capacity of our network from 2 Mbps to 40 Mbps in the last 10 years. 100 Mbps is projected to be available in the next 3 years.

- Beyond adding raw capacity to each serving group, we have been splitting the serving groups into smaller segments. The number of customers served from a fiber node (where our network converts from fiber to coaxial cable) has declined from an average of approximately 1,200 to less than 400 in the last 10 years, enabling us to deliver greater bandwidth to each end user.

- In addition, "above" the hybrid-fiber-coax network, we have been improving our baseline router and optical transmission facilities. In the past 10 years, capacity has grown from 45 Mbps to 10 Gbps per network node, an increase by a factor of more than 20. We are making initial investments to upgrade to 100 Gbps per channel in the near future.

The result of these investments has been dramatic – as noted above, a thousand-fold increase in average end user consumption per month, even as our data service subscriber base has increased from a very small figure in 2003 (when Bright House Networks obtained its systems), to approximately 1.5 million data services customers throughout our entire footprint today.

Despite the substantial cost of these investments, the price-performance of our broadband Internet access services has consistently improved over time. A simple but valid metric of price

[11] *See, e.g.*, http://www.cablelabs.com/cablemodem/primer/.

performance is the amount an end user pays per megabyte of information downloaded. This metric shows the same exponential increase as the technical performance metrics shown above:

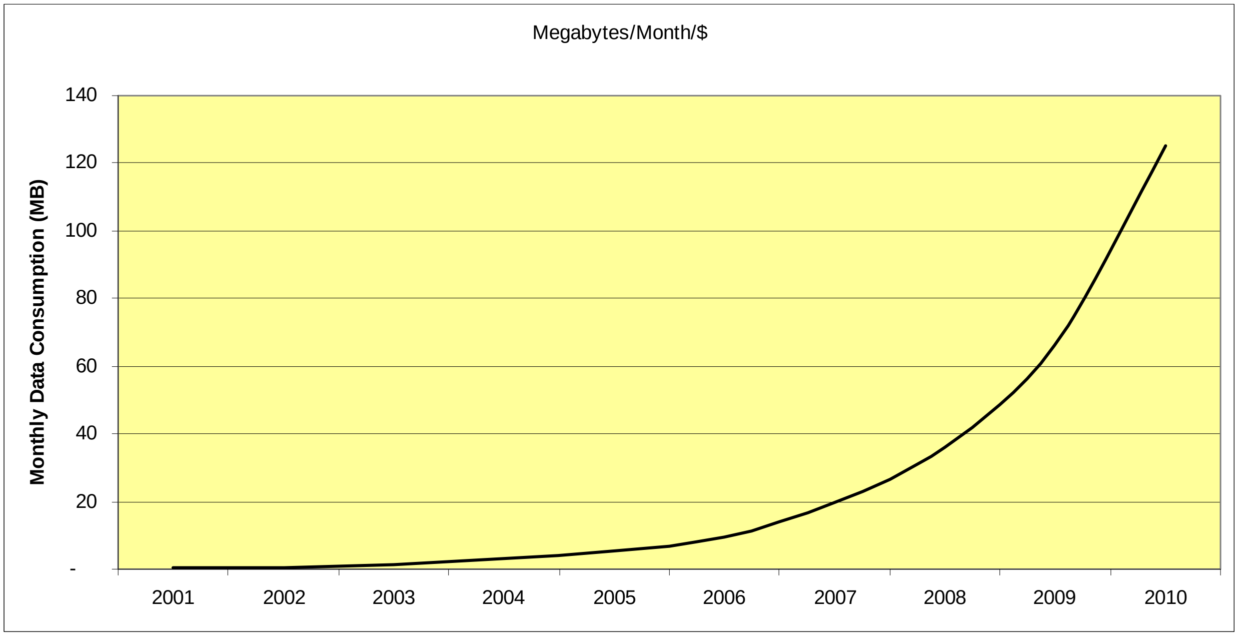

Bright House Networks' performance in improving its broadband Internet service – spending large amounts of money, with the express purpose of increasing the capacity of the service, while charging our end users less per unit of information – is not remotely consistent with the concern expressed in the Supplemental Notice, *viz.*, that retail network operators have some plan, desire, or incentive to constrain or restrict the capacity made available to broadband Internet access service. To the contrary, the evidence shows that we have been working hard to keep improving our network in order to keep up with end users' demands for more Internet capacity and functionality.

In this regard, while we are proud of our network and our efforts to improve it, the reason we have undertaken those efforts is no mystery – ***all*** retail network providers have done so to varying degrees, in response to consumer demand and competitive pressure. For example, in the Tampa Florida area, while we have been expanding our network capabilities, Verizon has been deploying, promoting, and upgrading its FiOS service. Were we to slow down our network upgrades, Verizon and others would immediately exploit that fact to take away customers.

In this regard, our competition is, of course, not limited to Verizon's landline service. Retail broadband Internet access service is also available from legacy "wireless" networks such as Verizon Wireless, AT&T Wireless, Sprint, and T-Mobile, which have all have been expanding – and touting – their data services.[12]

Moreover, our broadband Internet access service is increasingly purchased as part of a bundle, combined with our video and/or voice services.[13] If customers find our broadband Internet access service insufficient compared to that offered by Verizon, AT&T, or other competitors, we run the risk of losing the customer for ***all*** of the bundled services. In other words, in practical terms the risk we face if our broadband Internet access service is not competitive goes beyond the revenues from that service alone, and includes the revenues from the other services that are purchased along with the data service. This provides an added spur to ensure that our broadband Internet access service meets end user expectations.

All the various providers noted above are relevant in the retail market because end users overwhelming prefer devices that, if not literally "mobile" (like an iPhone), are at least entirely portable (like a laptop or notebook computer, or an iPad). The number of portable computing

---

[12] Furthermore, retail outlets ranging from McDonalds to Starbucks to Cosi to Panera Bread offer free WiFi to anyone in their establishments. Colleges and other institutions often provide free WiFi as well. As discussed more fully in Section IV, below, there is no reason today to distinguish among retail broadband Internet access networks based on what legacy service (if any) the network operator might originally have focused on.

[13] The Commission is well aware of the increased use of bundles in selling services to end users. *See, e.g.* FCC Wireline Competition Bureau, Industry Analysis & Technology Division, Local Telephone Competition: Status as of June 30, 2009 (released September 2010), at Figure 4 (showing overwhelming majority of VoIP services sold by ILECs and non-ILECs alike bundled with broadband Internet access service).

devices sold has eclipsed the number of non-portable devices (desktops).[14] These portable devices come with embedded WiFi and other wireless capability. With these devices in hand (literally), our base of retail end users can obtain broadband Internet access service at an ever-widening set of locations, for free. While we do not mean to suggest that free WiFi is a complete substitute for high bandwidth Internet access service in the home, the growing dominance of portable devices, the fact that most people are not at home most of the time, and the availability of broadband Internet access outside the home (whether free or not) affects the relative value to end users of the in-home services we offer. This is simply not a market environment in which we can sit on our capacity and have any hope of surviving, much less prospering.

In these circumstances, Bright House Networks submits that the notion that "the open Internet may wither as an open platform," as some have suggested, is a bit far-fetched.[15] To the contrary, all indications are that "the open Internet" is growing exponentially and will continue to do so.

B. Many Specialized Services Cannot Effectively Be Offered Over The Best-Efforts Internet.

The best-efforts Internet is a remarkable general-purpose communications network.[16] It is protean in its ability to handle everything from email, to interactive gaming, to voice

[14] "Laptop Sales Overtake Traditional Desktop PC Sales Over Holiday," http://www.switched.com/2009/01/09/laptops-overtake-traditional-desktop-pcs/ (Jan. 9, 2009); Mitchell, "Decline of the Desktop" http://www.computerworld.com/s/article/104856/Decline_of_the_Desktop (Sep. 26, 2005).

[15] Supplemental Notice at 2, *citing* comments of the Center for Democracy and Technology, Google, and others.

[16] The term "best-efforts Internet" is used to refer to, and highlight, the fact that the protocols of Internet routing and transmission provide no guarantees that the transmission of information will occur at any particular level of quality, or at all. *See* H. Newton, NEWTON'S TELECOM DICTIONARY, 19TH ED. (2003) at 106 (defining "best effort" as a "quality of service … class with no specified parameters and with no assurances that the traffic will be delivered across the network to the target device").

telephony, to video streaming. At first blush it seems that the Internet would be a good way to handle any communications function one could imagine.

In fact, however, the Internet has limitations that make it less-than-optimal, and in some cases completely inappropriate, as a means to handle certain services. As a result, whatever the Commission might do with regard to specialized services, it should make sure that its rules and policies do not constrain network operators' ability to make available to residential, business, and government end users, various services for which the best-efforts Internet is simply not a good delivery vehicle.

Although there are doubtless any number of technical matters that can affect the suitability of the Internet to handle particular services, broadly speaking, there are three areas of concern regarding the best-efforts Internet. First is that the Internet protocols are indeed "best efforts" oriented, with no assurance – and no way to provide assurance – that any particular level of quality will be available on an end-to-end basis (including the question of whether a particular transmission will go through at all).[17] Second, transmission of large amounts of data across the best-efforts Internet will inevitably encounter issues with latency, *i.e.*, intermittent or persistent lags in the time it takes for packets to make the end-to-end transit.[18] Third, the fact that the

---

[17] *See, e.g.,* Google Comments at 75 n.227 ("it is not obvious that a last-mile managed service accomplishes much of anything; only an end-to-end prioritized treatment could arguably be a demonstrable improvement over best-efforts open Internet access").

[18] This problem arises, among other things, from the combination of a multiple-router-hop path between two end points, combined with the operation of the Internet's "slow start" algorithm. Under that algorithm, one device sending packets to another across the Internet slows down the rate of transmission in response to actual or apparent congestion at any of the routers along the path between the two devices. *See* http://en.wikipedia.org/wiki/Slow-start. This means that even intermittent congestion at routers in a multiple-router path can cause the end-to-end transmission rate to vary significantly, which can lead to significant latency in transmission. Another source of latency is the "window" size associated with a service. *See* http://www.tcpipguide.com/free/t_TCPWindowSizeAdjustmentandFlowControl.htm. *See also* http://en.wikipedia.org/wiki/TCP_window_scale_option. While bandwidth is not irrelevant to latency, and insufficient bandwidth along one or more links may contribute to latency, problems relating

Internet is open to essentially anyone in the world raises the prospect that a network or service could be intentionally compromised or brought down by someone with malicious intent.[19] Any service that reasonably requires quality guarantees, or cannot tolerate substantial latency, or that must operate in a secure fashion, is not a good long-term candidate for delivery over the Internet. Network operators, therefore, must be free to offer these kinds of services on a "specialized" basis, without restriction.

Many services, in one or another of these respects, require more than the best-efforts Internet is capable of giving. For example, even very basic telemedicine applications – where a doctor is able to conduct a reasonably thorough visual examination of, and interview with, a distant patient – would require reliable, low-latency connections. On the "data" side of telemedicine, reliable connectivity between home-based health monitoring equipment and a medical provider's office could well require reliability levels high enough, and latency levels low enough, that special handling of the communications would be necessary. And, without question, advanced telemedicine services, where a doctor in one location would be actively controlling a health procedure being performed in another location, in real time, cannot reasonably be expected to rely on the best-efforts Internet.[20]

Similarly, real-time "telepresence" applications require a combination of high bandwidth and low latency that the Internet is not well-equipped to deliver. These problems are not severe enough to matter for free video-chat services like Skype. However, that level of quality is not adequate for professional business meetings or similar applications. As a result, while

---

to the slow-start algorithm and TCP window sizing can and do arise even when transmission bandwidth itself is not a constraint affecting a particular communication.

19 For these purposes this problem would include malware such as worms and viruses, as well as network-based actions such as botnet-driven distributed denial of service attacks, security breaches based on hacking into a supposedly secure password system, and others.

20 *See, e.g.*, CDT Comments at 47-48.

businesses are increasingly using telepresence services, they are typically carried via private (non-Internet) connections.[21] In this regard, there is every reason to think that telepresence services would be useful and valuable to individual end users as well as large businesses. Possible applications include everything from conducting parent-teacher conferences to allowing actors to practice their scenes, or musicians to practice challenging pieces of music together. These and other artistic endeavors require careful coordination – good timing – that cannot reliably be delivered using the best-efforts Internet.

Other applications are also subject to the reliability/latency problems described above. For example, as online gaming becomes more immersive and more graphics- and computationally intensive, even slight delays in relaying a particular player's actions back to the game servers and out to the other players can significantly degrade the game experience. Gaming, too, is a good candidate to be offered using a specialized service, rather than the best-efforts Internet.

From a more industrial perspective, sophisticated energy management services could well have communications needs that the Internet itself cannot meet. Simple online meter reading, or energy management that merely signals certain devices to shut down during periods of high demand, are not particularly challenging from a network perspective. However, more advanced "smart grid" approaches could not reasonably be trusted to the best-efforts Internet with its fluctuating, unpredictable latencies. For example, a service that would coordinate, in real time, the switching and transmission of grid-supplied power with many small bits of end-user-produced electricity (perhaps from solar cell arrays on the roofs of individual homes or

---

[21] *See* "Securing High-definition video conferencing and telepresence calls," available at: http://searchunifiedcommunications.techtarget.com/expert/KnowledgebaseAnswer/0,289625,sid186_gci1520953,00.html. As this article indicates, security is also a factor driving business and enterprise customers to use private networks rather than the best-efforts public Internet for telepresence services.

businesses) would require a level of reliability and coordination (low latency) that the best-efforts Internet cannot deliver.

Security is also a concern for essentially any service offered over the best-efforts Internet. In the industrial context in particular, we note recent press reports of the development of computer malware that appears to target industrial process control computers and software.[22] The evidently growing threat from such malware would caution against integrating power management and control systems – or any industrial systems with potential safety issues, such as gas or oil pipelines – with the Internet at large.

Other services provided using the same multifunction physical network that provides broadband Internet access service have reliability requirements that affect the use of capacity on the network as a whole. For example, a growing segment of our business is providing point-to-point private network connections for business and enterprise customers.[23] These data networks have stringent uptime and quality of service requirements that can only be met by dedicating significant network capacity to these services and devoting substantial resources to monitoring and maintaining their performance. These private data networks may involve transmission of customer information in IP format, but they do not directly connect to the Internet. (Of course, business customers can and do purchase high-speed Internet access from us as well, and may, if they choose, send their own data to locations on the Internet or receive data from the Internet.) In managing our integrated multifunction network, we must ensure that there is adequate

---

22 *See, e.g.*, E. Nakashima, *Stuxnet Malware is Blueprint for Computer Attacks on U.S.*, WASHINGTON POST, A Section (October 2, 2010), also available online at: http://www.washingtonpost.com/wp-dyn/content/article/2010/10/01/AR2010100106981.html.

23 For example, a healthcare company may want to establish a private data network linking several hospitals, medical office buildings, and other medical facilities so that medical imaging and other data can be transmitted where it is needed with minimal delays.

capacity to meet these customers' needs, along with those of our video, VoIP, and residential high bandwidth Internet access service customers.[24]

C. Under The Prevailing Intellectual Property Regime, Specialized Services Represent A Good Way – And Perhaps The Only Way – To Deliver Certain Highly Valuable Content To Individual End Users' Homes.

In addition to purely technical reasons to handle some services outside the confines of the best-efforts Internet, business considerations can also legitimately lead to the same result.

Many Internet sites are not oriented towards making money. For example, sites maintained by government entities, non-profit entities, or businesses may simply be designed to make information about their views, products or services available to a wide audience. For these entities, a presence on the Internet is not so much a money-making proposition (at least not in any direct way) as it is a cost of doing business.

Other Internet sites, obviously, are businesses focused on making money, but even these use different models. One common model is to present interesting information or applications to attract a large number of visitors, and then sell advertising. Another is to use an Internet site as a store, where people buy things; the money comes from the things being sold, not the content or capabilities of the site itself. Still another model, however, is to present certain information or capabilities to anyone visiting the site, for free, but to require payment to obtain access to other content or capabilities. In this latter case, what is being sold is the information on, or capabilities of, the site itself.

The point is obvious, but it bears emphasis that this latter model – selling content online – is only possible because the material itself is copyrighted. Copyright law gives the rights holder

---

[24] We note that some commenters, such as the Center for Democracy & Technology, explicitly recognize that these kinds of services should be permitted without regard to "neutrality rules." CDT Comments at 47.

the power to decide when, whether, and on what terms the material will be made available. Copyright holders have no obligation to make their material available over the Internet at all, or to make it available for free, or to make it available free of technical or legal restrictions on when or whether it can be copied or used.[25]

Again, these facts about American copyright law are obvious, but their impact on the question of specialized services does not appear to have been recognized in the Supplemental Notice. If some entity holds the rights to some particularly valuable content – say a recently released high-definition movie – that entity has no obligation to make that content available over the Internet – or any other distribution network – ***at all.*** A decision not to put content on the Internet may arise from technical concerns of the sort outlined above.[26] But it may be that the rights holder wants to use an entirely different – non-Internet – model for delivery of its content for business or economic reasons. In that case, the rights holder and a retail delivery network would have to agree on everything from the technical means by which the content will be delivered, to business questions regarding who will bill end users, what prices to charge, etc.[27]

This kind of arrangement would certainly be a "specialized service," but there is no conceivable reason to prevent rights holders and retail network operators from developing and implementing different technical and business models to deliver highly valuable content. Given

---

[25] To the contrary: as shown by the Digital Millennium Copyright Act ("DMCA"), the policy of the United States is to allow rights holders to establish technical means that limit access to their materials, with criminal and civil penalties applying to those who would circumvent those technical means. *See* 17 U.S.C. §1201 (banning circumvention of technical means of securing copyrighted material); § 1203 (civil penalties for violations of § 1201); § 1204 (criminal penalties for violations of § 1201).

[26] High-definition video, for example, would suffer greatly from problems of excessive latency.

[27] For example, one can imagine an owner of valuable content testing a model that might deliver the content to a local retail network operator via the Internet in heavily encrypted form, but that would require the retail network operator to store it, decrypt it, and stream it to end users in a particular format over bandwidth or channels provided on a "specialized" basis. In a related context the FCC found that it is in the public interest to allow innovative business arrangements when it eased the prohibition on selectable output control to allow consumers to obtain early-release films. *See Motion Picture Ass'n of America, Petition for Expedited Special Relief*, Memorandum Opinion & Order, 25 FCC Rcd 4799 (2010).

that it is the rights holders – not retail networks, and not the Commission – that are legally entitled to decide whether and on what terms to make their content available, the only effect of restricting the ability of retail networks to establish specialized services arrangements in this context is that certain valuable and interesting content will not be made available to end users in their homes at all. This would clearly not serve the public interest. This is another reason that the Commission should not impose restrictions with respect to specialized services that a retail network operator might offer.[28]

* * * * *

The examples of specialized services discussed above show that skepticism about the need for, or legitimacy of, such services, is unfounded. There are now, and likely increasingly will be, services that end users need and that a network operator will be able to provide – using the same physical network used for Internet access – that are simply not appropriate for delivery over the Internet itself. Network operators will, for the foreseeable future, offer both broadband Internet and non-Internet services over the same physical facilities – which, as noted above, is

[28] We are of course aware of the concerns expressed by some thinkers that copyright law grants too much control over "cultural" production to rights holders, and that rights holders should either choose to, or should be required to, make their materials available on more open terms – including on the Internet. *See generally, e.g.,* Y. Benkler, THE WEALTH OF NETWORKS (2006); L. Hyde, COMMON AS AIR: REVOLUTION, ART & OWNERSHIP (2010); L. Lessig, CODE AND OTHER LAWS OF CYBERSPACE (1999); L. Lessig, THE FUTURE OF IDEAS (2001); J. Littman, DIGITAL COPYRIGHT (2001); J. Zittrain, THE FUTURE OF THE INTERNET AND HOW TO STOP IT (2008). Whatever the force of these arguments, they do not change the fact that the law today grants very extensive rights to copyright holders, and retail network operators are not free to ignore those rights. For example, in addition to the basic fact that we cannot offer copyrighted content without arrangements with the rights holder, we incur significant expenses responding to demands from rights holders under the DMCA that we "take down" allegedly copyrighted material posted by our end users. Indeed, all participants in the Internet ecosystem incur expenses, and to some extent have to design their systems, in order to accommodate the intellectual property rights of content owners. Notwithstanding this fact, it seems clear that the Commission cannot set rules – regarding specialized services or otherwise – with an eye towards tilting intellectual property rights in one direction or another.

what they have always done.[29] In addition, it is content owners, not retail network providers, who have the right to control whether their content will be made available by means of the best-efforts Internet. If rights holders are unwilling to make their available via the Internet, but are willing to make it available by means of specialized service arrangements, the public interest is served, not harmed, by permitting those arrangements. And the fact that such content may be available via a specialized service does not create any incentive on retail network operators to degrade the functionality of their broadband Internet access services – for which there is substantial and growing independent demand – or to fail to continue to increase the capacity available for such services.

For all these reasons, Bright House Networks submits that concerns that network operators will constrict or interfere with broadband Internet access as a result of also offering other, non-Internet services are not merely speculative; such concerns are affirmatively unrealistic. Bright House Networks' end users want broadband Internet access, so we will continue to work to provide it to them. Our end users also want other services (which today include, for example, our traditional cable service and our VoIP service, and in the future will certainly include various specialized services of the type identified above), and we will continue to provide those services as well. At any given moment our network will be configured to

---

[29] The fact that the traditional Internet is not the best way to deliver all services should not be viewed as particularly remarkable. As technology and consumer demand change and evolve, the optimal serving arrangements change as well. *See* Section III.B.3, *infra*. For example, the traditional web browser-based interface is turning out not to be the best way to deliver all Internet-based services. For many years the traditional world-wide web, involving sites scripted in hypertext markup language (html), and accessed by a browser using the hypertext transport protocol (http), seemed to be essentially the only way to offer services over the Internet. This is one reason the so-called "browser wars" seemed so important in the 1990s. Increasingly, however, it appears that the traditional world-wide web does not offer enough ability to manage a user's experience for many developers, who instead write customized, non-web-based applications. *See* C. Anderson & M. Wolff, "The Web Is Dead: Long Live The Internet," WIRED 18.09 (September 2010) at 118. Just as the Internet community is coming to understand that not all services must or should be offered via the web, so too the networking community understands that not all services should be offered via the Internet itself.

allocate its then-current capacity among various services – Internet, cable, VoIP, and others – in a manner that reasonably balances demand for all the services our end users. But if anything is clear from the last decade of experience, it is that over time each of the different services – and, specifically, broadband Internet access – will be allocated more capacity, not less.

III. THE OTHER STATED CONCERNS REGARDING SPECIALIZED SERVICES ARE ILL-DEFINED AND, IN ANY CASE, MISUNDERSTAND THE INTERNET ECOSYSTEM AND RETAIL NETWORK OPERATORS' ROLE IN IT.

The Supplemental Notice indicates two additional areas of concern regarding specialized services: (1) that a particular service might be offered on a specialized basis rather than via the "open Internet" in order to supposedly "bypass open Internet protections,"[30] or (2) that a service might be offered on a specialized basis in order to engage in some unspecified "anti-competitive conduct."[31] Unlike the concern that network operators will cut back on broadband Internet capacity – which, while unrealistic (as described above), is at least stated plainly enough – these concerns are not spelled out with any clarity, making a precise response difficult.[32] We do, however, offer the points below for the Commission's consideration.

A. As Long As Broadband Internet Access Is Offered In An Open And Nondiscriminatory Manner, No Valid Policy Goals Are Served By Restricting Specialized Services.

The Supplemental Notice assumes that the Commission will establish enforceable obligations regarding openness and nondiscrimination regarding broadband Internet access, but then questions whether specialized services would undermine those obligations. We do not believe such regulations are required in the first place, because broadband Internet service is

[30] Supplemental Notice at 2.

[31] *Id.* at 3.

[32] The comments the Commission cites in raising these concerns are not particularly helpful in explaining them. For example, the Commission cites Netflix Comments at 9-10 as part of the record basis for both of these concerns, but Netflix simply states, generally, that network operators should not use specialized services to (in some completely unspecified way) "end run the openness and nondiscrimination rules" or to "adversely impact the provision of a robust, open Internet access service."

already offered in an open and nondiscriminatory manner. This is the result of the competitive pressures noted above, and also of the fact that the value of the service arises from its open and nondiscriminatory nature.[33] As long as the capacity devoted to open and nondiscriminatory broadband Internet access service is large, and growing – which it is – the fact that the network operator also offers other services using additional network capacity is simply irrelevant to any legitimate public policy concern.

Indeed, engaging in regulatory action based on these ill-defined concerns is an invitation to make policy that is anti-consumer, anti-competitive, and anti-innovation. For example, some commenters appear to think that existing cable services should be considered specialized services under new Commission rules.[34] One commenter evidently sees regulation of specialized services as a way to insulate itself from competition, stating that a network operator might obtain regulatory advantages by "position[ing] their voice and video services as 'Managed' or 'Specialized,' exempting them from the Commission's proposed principles."[35] Of course, cable operators such as Bright House Networks ***already*** offer voice telephony (VoIP) and video services (cable) on the same network used to offer Internet access. Those services are subject to regulatory regimes that are entirely distinct from whatever openness or nondiscrimination obligations the Commission may ultimately impose with respect to broadband Internet access service.[36]

---

33 That is, people want to buy broadband Internet access precisely because of the wide range of services easily and equally available using it.

34 *See* Netflix Comments at 9 (referring to network operators' "legacy 'managed services'"); Vonage Comments at 27 (expressing concern that a network operator might obtain regulatory advantages by "position[ing] their voice and video services as "Managed" or "Specialized," exempting them from the Commission's proposed principles")

35 Vonage Comments at 27, *supra.*

36 The Communications Act recognizes that entities will provide multiple services over the same physical network, and that those services will be subject to different regulatory rules, depending on their

The troubling implication of Vonage's comment is that, because some forms of video and voice services are available on the best-efforts Internet, cable operators and other network providers should banned from offering their VoIP or cable services at any level of quality higher than that available with the Internet-based services. In other words, Vonage is asking the Commission, in the name of "nondiscrimination," to prevent its competitors from offering superior service. We trust that the Commission will not entertain suggestions that existing services must be "dumbed down" – or efforts to improve them curtailed – because an Internet-based alternative is or becomes available.

B. The Internet Ecosystem Works Cooperatively To Deal With Issues Such As Those Presented By Specialized Services.

No one member of the Internet ecosystem produces a finished service of value to end users. Instead, the value of the Internet arises from the joint production, by the ecosystem as a whole, of finished services on an end-to-end basis.[37] As a result, all participants in the Internet ecosystem have a mutual interest in delivering end users the services and applications that they want, in the most efficient way possible. This does not create a situation in which members of the ecosystem exploit each other. It creates a situation in which the ability of any member of the ecosystem to deliver value to end users requires ongoing coordination and cooperation. The

---

respective characteristics. For example, an entity that is a "telecommunications carrier" with respect to some of its services may not be regulated as a carrier with respect to its other, non-telecommunications services. *See* 47 U.S.C. § 153(44) ("A telecommunications carrier shall be treated as a common carrier under this Act ***only*** to the extent that it is engaged in providing telecommunications services"). *See also* 47 U.S.C. § 541(c) ("Any cable system shall not be subject to regulation as a common carrier or utility by reason of providing any cable service"). Under the Act, therefore, there can be no claim that the provision of broadband Internet access using a particular physical distribution network subjects all the other services on that network, or the physical plant itself, to the same kinds of regulation that might apply to Internet access service, or to some form of common carrier regulation. Any Commission action with respect to specialized services must be consistent with these statutory principles.

[37] For example, no matter how fast and capacious Bright House's own network might be, without access to all the content and applications on the Internet – such as Google, Facebook, YouTube, Yahoo!, Windows Live, Wikipedia, Blogger.com, Twitter, MSN, and Amazon.com – it is hard to see why end users would want to buy our service in the first place.

result is a "commons," a system in which members' actions are governed not by the raw economic incentives of the market, nor by any outside regulatory obligations, but, instead, by the internally-generated obligations of the system itself. In such an environment, there is no reason to believe that significant abuses of one member of the commons by others – the essence of the "bypass the open Internet" and "anticompetitive conduct" concerns stated in the Supplemental Notice – will occur or persist. Moreover, the sheer complexity of this commonly produced resource requires that issues be sorted out by the community itself, not by external regulatory commands.

1. The Physical Infrastructure Of The Internet Is Produced As A Commons, With Its Own Internal Processes For Resolving Resource Management And Allocation Problems.

It is commonplace to note that "the Internet" is not regulated. This is true enough in the sense that the Commission has not historically had rules and requirements regarding connections between retail networks and the backbone, peering, etc. It is also frequently stated that the Internet, being unregulated, is controlled by "market forces." That formulation, however, is overly simplistic. First, many important arrangements regarding the architecture of the physical Internet – notably the peering and other interconnection arrangements by which the vast majority of Internet traffic is exchanged – are handled on a settlement-free basis – *i.e.*, without a price, which makes reference to "market forces" a bit odd. Second, the underlying technical rules for the Internet – the standards and protocols embodied in Requests for Comment (or "RFCs") are developed by consensus, via groups such as the Internet Engineering Task Force ("IETF") and the Internet Architecture Board ("IAB"), not in any market-based system. Third, certain critical resources needed to make the Internet work – notably, Internet addresses and domain names – are administered by commonly managed entities such as the Internet Corporation for Assigned Names and Numbers ("ICANN") and (in the U.S.) the American Registry of Internet Numbers

("ARIN"), not by any market-based allocation mechanism. Fourth, as noted above, as a practical matter the members of the Internet ecosystem, while legally and operationally separate entities, are not really "competitors" in any standard economic sense with respect to the provision of the end-to-end capacities of the Internet infrastructure, because they are all jointly involved in providing that infrastructure, which is then used, as a whole, to provide services that end users want to buy.

For many years economic and policy thinking assumed that there were only two stable organizational models for allocating resources: free markets or command-and-control regulatory regimes. More recent economic analysis, however, has focused on a third way of looking at resource allocation – the concept of a "commons."[38] A commons is a jointly used productive resource that must be managed in order to preserve and promote its value over time. These joint decisions regarding managing the common pooled resource are not, and cannot be, based on simple market forces – allowing unbridled market forces to operate in this kind of situation produces the well-known "tragedy of the commons."[39]

---

38 *See* E. Ostrom, GOVERNING THE COMMONS (1990) at 1 ("communities … have relied on institutions resembling neither the state nor the market to govern some resource systems with reasonable degrees of success over long periods of time"). Some draw a distinction between a "limited-access common property regime" that is open only to a limited group and treated as property as against non-members, and a commons that is literally and fully open to everyone,. *See* Benkler, *op. cit.* at 61. As Professor Ostrom points out, however, if a resource may be used by everyone without depleting it, the resource is a non-rival "public good," not really a commons at all. Ostrom, *op cit.,* at 32. The point of identifying a "commons" is that, despite its "common" aspects, it remains a resource that requires management in a way that non-rival "public goods" do not. For this reason, analysis treating "culture" or "intellectual property" – or other non-rival goods – as a commons, *see, e.g.*, Benkler, *op. cit.*, Hyde, *op. cit.*, is not directly applicable to the Internet infrastructure itself, which is a physical resource with finite capacity (at any given time) that must be managed in order to function properly.

39 That is, without management, some with access to the common resource could over-use it to such an extent that it would become less valuable, or perhaps valueless, to everyone else. Part of the impetus for Professor Ostrom's work was the observation that this "tragedy of the commons" is, in fact, frequently avoided in practice. *See* E. Ostrom, *op cit.*, at 21. By way of example on the Internet, it was recently reported that more than of 90% of the emails sent over the Internet are not actual, legitimate communications between end users, but instead are spams. *See* http://news.cnet.com/8301-1009_3-

The members of a commons each have their own economic interests, and may well sell the output of the common productive resource to outsiders. But their pursuit of those interests will be constrained both by their awareness of the mutual dependency of all members of the commons, as well as by the rules, customs, and institutions established to allow the members to pursue their own interests while still preserving the viability of the shared resource on which they all depend.[40]

The "commons" model clearly fits the way in which the physical Internet infrastructure is produced. First, as noted, no one entity produces an end-to-end valuable resource; instead, the various members of the commons each produce something that is only valuable in conjunction with what the others produce.[41] Second, if members of the commons pursued only their own short-term interest without regard for the commons as a whole, a "tragedy of the commons" arises.[42] Third, the commonly produced and managed resource exhibits positive externalities, so

---

10249172-83.html. Without active management of email systems throughout the Internet, spam would simply overwhelm the ability to make productive use of normal email.

40 "When multiple [members of a commons] are dependent on a given [common resource] as a source of economic activity, they are jointly affected by almost everything they do. Each individual must take into account the choices of others when assessing personal choices. … The key fact of life for [members of a commons] is that they are tied together in a lattice of interdependence so long as they continue to share a single [common resource]." Ostrom, *op. cit.* at 38.

41 For example, no matter how fast and capable an end-user-facing network might be, no one will want to attach to it if it does not offer access to the wide variety of applications, content and services available on the wider Internet. No matter how interesting or useful a provider's content or application might be, no one will be able to obtain it if it is not connected to the wider Internet and its millions if not billions of connected end users. A backbone network is only useful if there are users and content/application providers who need to connect with each other. A data center is only useful if both backbones and retail networks exist to link the hosted data to end users who want it. Etc.

42 *See* note 39 *supra* (example of spam). If various members of the Internet community did not take active steps to manage spam and prevent it from reaching end users, the value of email as an effective and efficient communications system would be destroyed, as end users were forced to spend enormous amounts of time deleting unwanted emails. *Cf.* Ostrom, *op. cit.* at 23 (noting problems that can arise when a limited-access common property resource is converted, without regard for resource management, into an open-access resource; *see* note 38, *supra*).

each member of the commons benefits, directly or indirectly, from the activities of the others.[43] It is precisely because of these economic characteristics that purely market-based decision making is not viable. This had led to the establishment and persistence of the unique non-market governance and resource allocation systems noted above (ICANN, ARIN, the IETF's RFC process, etc.).[44]

A commons operates with a very different set of incentives than exist in a market. Notably, the members of the commons understand that it is the coordinated efforts of all of them, not any one player's individual action, that produces the value that all of them need to survive economically.[45] Moreover, and especially due to the highly technical and complicated nature of the Internet, each member of the commons will have repeated opportunities to provide assistance to others, but also knows that it will need assistance from others in the future. At the same time, it is impossible to specify in advance (and, therefore, impossible to monetize) precisely what assistance will be provided or received.[46] There is, and can be, no market that sets prices for

---

[43] For example, when some group puts up a new web site, that adds value to everyone's Internet service, even though most of the time most people will not access it, or even know it is there. From one perspective, this is a special instance of "Metcalfe's Law," which state that the value of a communications network increases with the square of the number of nodes (users) on the network. *See* http://en.wikipedia.org/wiki/Metcalfe's_law.

[44] In this regard, note that until the Internet was opened up for commercial use in the early 1990s, the physical infrastructure of the Internet was owned and managed by government and educational institutions. *See* http://en.wikipedia.org/wiki/History_of_the_internet. It seems that the long history of cooperative, non-commercial management of the Internet prior to the introduction of commercial activity created a community of network engineers and operators who understood the need to manage the Internet resource cooperatively in a way that may well not have occurred if pure economic "market" forces had been fully operational from the outset.

[45] *See* note 40, *supra.*

[46] A typical example of these kinds of accommodations within the Internet commons is settlement-free peering, where different networks recognize that they both benefit from connecting to exchange traffic at one or more points, and do so at no charge for the traffic exchange. This general phenomenon is well-described, in a slightly different context, by Professor Benkler. Y. Benkler, *op. cit.* at 106-116. His analysis explains how, in effect, people's ability to keep an informal track of favors given and favors received makes it possible for mutually beneficial transactions to occur even in the face of substantial uncertainty when any the specific "payment" (in the form of future favors) will be received, or even what

these mutually beneficial actions. Instead, they are, of necessity, worked out informally and on a case-by-case basis among the members of the commons.

Recognizing that the physical infrastructure of the Internet (broadly construed) is provided by a commons has two key implications here. First, the mutual dependence of the members of the commons on each other means that it is not accurate to assume that any member of the commons will act in its own short-term interest to the detriment of other members. Second, there is no reason for the Commission to be concerned – in advance and based on speculation – that the Internet community will be unable to resolve any problems that might arise in connection with the offering of specialized services. To the contrary, as described below, there is every reason to expect the community to deal with such issues without outside intervention.[47]

---

form it will take. From another perspective, this situation is an example of a game in which the players are able to reward and punish each other, but in which they all know that the game will be repeated many times. In such situations, empirical work in game theory and behavioral economics shows that the players will tend to treat each other fairly, presumably in order to encourage fair treatment of them by others in the future. *See.*, *e.g.*, http://en.wikipedia.org/wiki/Prisoner's_dilemma; *see also* E. Ostrom, *op. cit.* at 6 (discussing M. Olson's point that the ability of a group to effectively organize is enhanced when members can observe each other's actions).

[47] Thus, the status of the physical Internet as a commons provides a sound basis for the Commission to conclude that, with respect to matters such as reasonable network management practices, specialized services, etc., it is entirely reasonable to rely on the commons' own institutions for reaching pragmatic consensus – such as the IETF, the IAB, etc. – to resolve disputes in the first instance. *See* E. Ostrom, *op. cit.* at 17. Regulatory intervention would be a last resort, if needed, to enforce the "common law" of the Internet. *Cf.* C. Rose, "The Comedy of the Commons: Custom, Commerce and Inherently Public Property," 53 U. CHI. L. REV. 711, 742, 746 (1986) (noting that the "custom of the country" ***was*** the "common law" that British courts would enforce). In this regard, the recent establishment of the Broadband Internet Technical Advisory Group (BITAG) reflects the type of action that is needed to develop solutions, based on input from knowledgeable people who will be able to reasonably assess not just the interests of whatever organization they might formally represent, but also the interests of the shared resource – the Internet – as a whole.

2. The Sheer Complexity Of The Internet Makes It Necessary To Rely On The Commons To Sort Out Issues Relating To Specialized Services.

One reason that it is so important to allow the Internet commons to function in its normal, informal way – including with regard to specialized services – is that the issues involved can be extremely complicated, both technically and economically. Case-by-case discussion and experimentation are typically needed to resolve any but the most trivial problems.[48]

This is true because any communications functionality involves some combination of three different factors – bandwidth (transport), storage (in the network, at the end points, or both), and computing power (again, in the network and/or at the end points).[49] These capabilities are, to some degree, substitutes for each other.[50] They are also embedded, to various degrees, in a wide variety of equipment and software, including end user devices, severs, routers, optical networking equipment, and so on. Moreover, the relative costs of the three factors are constantly changing as new transmission, switching, routing, and storage technology is developed and new software is written. As a result, the best way to attack some problem of communications

---

48 Ostrom, *op. cit.* at 17 (one cannot reasonably expect an external authority to have the data needed to properly direct resource allocation and other decisions by members of a commons); *id.* at 18 (a commons can work with members of the commons monitoring each others' behavior with regard to the common resource).

49 Benkler, *op. cit.*, at 2-3, 32.

50 For example, consider the task of delivering a movie to someone's home. One way is to use almost entirely bandwidth – with no storage and only minimal computation. That would be analog broadcast TV. Another way is to use almost entirely storage – with no bandwidth and, again, minimal computation. That would be the Blockbuster and (old) Netflix model of physically delivering DVDs (pure storage) to people's homes. Another way is to use a mixture of bandwidth, storage, and computation – deliver the movie via a network to the home, possibly in compressed or encrypted form, temporarily store it there, and then apply computer power to retrieve, decompress, and decrypt the file. That would be the (new) Netflix model of online downloads of movies.

networking – that is, how to most efficiently and reliably get data from its source to its destination – is a constantly changing optimization problem in network engineering.[51]

This applies fully to any newly-deployed specialized services, but particularly to those that are being made available on a specialized basis because some technical requirement of the service makes the best-efforts Internet an inappropriate delivery vehicle. It will almost certainly be impossible to determine in advance, or in the abstract, what particular technical configuration

---

[51] The complexity of the overall Internet ecosystem can sometimes be lost in the typical focus on "bandwidth" or "download speed" as a key indicator of quality – particularly for retail marketing purposes. For an end user to obtain and use almost any Internet service or functionality, however, depends on a complex interaction among at least the following:

- The capabilities of the computers and other Internet-enabled devices under the end user's control. This includes maximum connection speed, application buffer size (which can vary by application), the speed with which an application can input and process data on a particular computer (which will vary by computer and by application), and the degree to which a particular application will download and work with locally stored information rather than requiring more or less continuous, real-time downloads from a remote location.
- The bandwidth available on the "last mile" network connecting the customer to the Internet.
- The degree to which the Internet information or application the end user seeks to access is located "near to" or "far from" the customer, in terms of network topology (number of router hops required to reach the "far" end). A customer will have a very different experience in accessing data or an application that is "locally" available by virtue of the application provider's use of a service such as Akamai or Limelight, as compared to an application provider that uses a server that is only available at the "far" end of a multiple-router path through the Internet backbone.
- The degree of congestion – even transient congestion – on whatever set of routers the end user's inbound and outbound data packets might have to traverse.
- The capabilities of – and degree of congestion facing – the servers at the "far end" of the path between the customer and the application provider. A customer with a gigabit-per-second link into the Internet, for example, will still experience delays in downloading a web page or accessing an application that resides on a server connected to the Internet by means of a 512 Kbps DSL line, or a server that can only internally process data at a rate of 1 or 2 megabits/second. And delays may still occur between our hypothetical gigabit-per-second customer and a screamingly fast server, with its own multi-gigabit connectivity to the Internet, if one million, or ten million, high-bandwidth end users all want to access the same data or application at the same time.

The fact that there are so many "moving parts" involved in providing even simple Internet functions like downloading a web site is both a blessing and a curse. It is a curse because it can be very difficult to know the specific cause of some intermittent problem. *See, e.g.*, CDT Comments at 47 n.163; Google Comments at 75 n.227. But over the longer term, it is a blessing because there are many different ways for the multiple players in the Internet ecosystem to solve any real, persistent problem that might arise.

will be best suited to meet the service's actual technical needs. There may be different ways to attack the problem, and different retail network providers may want to take different approaches. The only effect that Commission rules or requirements could have in such a situation would be to add a layer of uncertainly and confusion onto the Internet community's already difficult job of trying to figure out some technical way to make things work.

3. Recent Experience Shows That The Internet Commons Will Naturally Treat Traffic With Special Requirements Differently, Without Treating Anyone Unfairly.

A good example of the Internet commons functioning to deal with a challenge regarding the requirements of particular traffic is the rise of semi-private "content delivery networks" over the last several years. In the model of the Internet as it existed in the late 1990s, someone with content to deliver would select an ISP to host the content. The ISP would connect upstream to the Internet backbone, and, thereby, to the universe of end users. An end user request for the content would traverse the backbone to the content provider's ISP. The content would then make its way back over the same basic route. This works if the content being delivered is a static file (say, a technical paper). But for the reasons described above, it does not work very well for applications, such as streaming video, where the packets need to arrive at the end user in proper order and without gaps or delays. The fact that the end user would experience problems with the delivery of video or similar content, was not the "fault" of anyone in particular – it was a function of the design of the Internet and its routing algorithms.[52]

The response of the Internet ecosystem to this problem has been the creation of distributed content delivery networks that bring popular content – including a great deal of video – much "closer" to the end users requesting it. For example, companies such as Akamai and

52 *See* note 18, *supra.*

Limelight have deployed servers to many locations around the Internet, and maintain copies of their customers' content and applications on servers close to end users. The retail networks recognize the value of this service and frequently, if not uniformly, offer settlement-free peering arrangements to these companies. Firms with very large data delivery needs of their own, such as Google's YouTube service, self-provision these distributed delivery networks and are typically accorded the same settlement-free peering arrangements. The result is that the technical problem of latency is solved by, in effect, bypassing the backbone. With the ongoing growth in video and related high-bandwidth content, a significant and ever-increasing fraction of Internet traffic traverses these non-backbone content delivery networks, rather than the backbone itself.[53]

The development of these backbone-bypass mechanisms is remarkable, and relevant to this proceeding, for two reasons. First, it shows that specialized services for Internet-delivered content in effect already exist, with no harm to the Internet as a whole. End users seeking to obtain content from a provider that does not have its own delivery network, and that does not use a company like Akamai or Limelight, may well experience delays in downloading. But no one seriously suggests, for example, that Akamai or Limelight be run out of business on the grounds that they give privileged status to some content over other content.[54]

---

[53] Labovitz, *et al.*, ATLAS Internet Observatory 2009 Annual Report, available on-line at http://www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf. As of the date of the study, approximately 10% of all Internet traffic – a huge proportion – originated from "pure play" CDNs such as Akamai, Limelight, etc. *See id.* at slide 15. Adding Google traffic (including YouTube) and other "private" CDN traffic such as Facebook's, results in an even greater proportion of traffic attributable to CDNs. Google alone, for example, now accounts for 6% of all Internet traffic globally. *Id.* at slide 18. This has resulted in a tremendous shift in the distribution of Internet traffic, so that in 2009 only 150 ASNs (autonomous system numbers) account for more than 50% of all Internet traffic. In contrast, only two years earlier, in 2007, thousands of ASNs were required to reach that level. *See id.* at slide 14.

[54] In this regard, just as retail networks like Bright House compete with each other to obtain end user subscribers and provide them with high bandwidth Internet access, so too do content delivery optimization services like Akamai, Limelight, Brightcove and others compete with each other to provide content and application providers with optimized (backbone-bypassing) delivery of their data to end

Second, this development illustrates the ability of the Internet community to deal with problems of this sort without regulatory intervention – because everyone benefits from having the problem solved. The community figured out a way to give content that needed special handling, the treatment it needed.[55]

In the specific context of specialized services, the experience with content delivery networks also illustrates that the Commission should be open to solutions that take certain traffic and certain services off of all or part of the traditional best-efforts Internet. Doing so may be necessary to meet the technical requirements of the service (as described above). Or it may be necessary in order to avoid imposing large and potentially unnecessary investments on other members of the Internet community. Or it may be necessary to test a particular business model.[56] Or it may make sense as a way to make use of a new approach to routing traffic.[57] Or it may make sense as an experiment to see whether a particular technical or business approach might work. There is no reason impose regulatory requirements on specialized services that would deprive the Internet ecosystem of the ability to try any of these approaches.

---

users. The existence of this competition for the business of individuals and entities that are not part of the Internet infrastructure commons in no way detracts from the status of the producers of that infrastructure ***as*** a commons, any more than the fact that fisherman compete to sell their catch means that their management of the common fishery is not a commons.

[55] The NANOG report cited above describes the internal, commons-driven reconfiguration of the Internet infrastructure as creating a "new core of interconnected content and consumer networks" that has resulted in "dramatic improvements in capacity and performance." Labovitz *et al.*, *op. cit.* at slide 17.

[56] *See* discussion in Section II.C, *supra.*

[57] For example, one of the "fathers" of the Internet, researcher Van Jacobson, has suggested that the basic approach of assigning Internet addresses to specific ***computers*** has become obsolete, given that today a typical end user is completely uninterested in which specific computer might hold the information in the end user seeks. As a result, Dr. Jacobson has proposed a new addressing and routing scheme that is based directly on the information end users seek. *See* V. Jacobson *et al., Networking Named Content,* CoNEXT '09, available online at, *e.g.*, http://www.cl.cam.ac.uk/teaching/0910//R02/papers/ccn2.pdf. This proposal is consistent with the observation by Labovitz, *et al.,* that the Internet is in "transition from focus on connectivity to [focus on] content." Labovitz, *et al., op. cit.* at slide 26. Bright House obviously has no idea whether Dr. Jacobson's new addressing and routing scheme will ultimately be adopted, but it serves as an example of the complexity of the factors that can go into devising an optimal and long-term cost-effective solution to often subtle networking challenges.

* * * * *

The discussion above shows that the Commission should not impose regulations restricting the ability of retail networks or other members of the Internet community to develop specialized means of delivering particular services to end users. The Internet commons has both formal and informal means of sorting out the technical and economic issues associated with such services, and, in any event, the issues they pose – and the solutions to those issues – are so technically complex that it would be a mistake to try to deal with them through *a priori* externally imposed regulations in any case.

IV. ANY REGULATORY REGIME ESTABLISHED TO ENSURE "OPENNESS" OR "NEUTRALITY" SHOULD APPLY TO ALL BROADBAND INTERNET ACCESS SERVICES, REGARDLESS OF TECHNOLOGY.

The Supplemental Notice also asks for additional comment on issues regarding the application of open Internet principles to wireless networks, focusing on questions of transparency in wireless offerings and capabilities of wireless devices and applications that run on them.[58] With regard to those issues, the Commission should not treat the networks operated by CMRS licensees differently from any other retail Internet access network.

As technology has continued to develop, any historical distinction between "wireless" and "wired" networks is entirely artificial, particularly, but by no means exclusively, in the realm of broadband Internet access services. From the end user's perspective, essentially every network that might be used for this purpose has a wireless component. In the home, for example, most supposedly "wired" broadband providers either automatically provide a wireless router for the end user or at least fully support such applications. Indeed, home wireless routers are

[58] Supplemental Notice at 4-6

standard consumer items, available at any number of retail outlets such as Staples,[59] Office Max,[60] and Radio Shack.[61] Recognizing the prevalence of WiFi-enabled devices, CMRS-based providers offer devices to allow customers to establish a WiFi "hotspot" anywhere their networks offer coverage.[62] And, of course, essentially all consumer devices with the ability to access the Internet at all – such as laptop computers, notebook and "netbook" computers, and newer devices such as iPads and iPod Touches – are WiFi-enabled. Indeed, mobile "phones" – such as iPhones and Android phones – are increasingly WiFi-enabled as well, so that a subscriber to a traditional CMRS service can use their device to send and receive VoIP calls over a local WiFi network rather than over the CMRS network.[63]

WiFi networks, of course, are only "wireless" within the relatively small local area covered by a given WiFi hotspot. From that point on, the data sent and received on a WiFi network is typically handled via a set of landline connections, such as a telephone company DSL line or a cable system, or a private line. For this reason, no one seriously suggests that the Commission's policies regarding broadband Internet access should not be applied to a WiFi-enabled network – which would exempt essentially ***all*** broadband Internet access from those policies. Yet the situation is the same for a CMRS-based data network: the service to the end user is only "wireless" from the end user's device to the nearest antenna/cell site.[64] From that

---

59 *See* http://www.staples.com/Wireless-G-Routers-Access-Points-Wireless-G/cat_CL142367.

60 *See* http://www.officemax.com/technology/networking-equipment/wireless-networking.

61 *See* http://www.radioshack.com/family/index.jsp?categoryId=2032370&s=null.

62 *See, e.g.,* http://www.verizonwireless.com/b2c/mobilebroadband/?page=products_mifi.

63 *See, e.g.,* http://www.slashgear.com/skype-for-android-released-free-3gwifi-voip-and-im-05105922/; *see also* http://icall.com/iphone/. This latter application claims to be able to switch a call in mid-stream from the cellular network to the WiFi network, thereby saving cellular usage charges.

64 As noted above, however, devices also exist to create a WiFi hotspot tethered to a CMRS data network. In that case there are two wireless "hops" rather than just one – from the WiFi-enabled device to the CMRS/WiFi hotspot; and from that hotspot to the nearest CMRS network antenna.

point on, the data sent and received on a CMRS-based "wireless" network is typically handled by means of landline connections, such as ILEC private line services.[65]

In these circumstances, if "wireless" service is to be given some special treatment under the Commission's policies for broadband Internet access, there is no reason to afford different treatment to networks that serve end users wirelessly via WiFi technology as opposed to wirelessly via CMRS technology. To the contrary, as just discussed, any distinction between WiFi and CMRS is blurring to the point of vanishing, as far as the end user is concerned.

Furthermore, as noted above, because end users are relying more and more on wireless devices for Internet access, there is no public interest basis to exclude wireless networks from any aspect of the Commission's policies regarding open Internet access – including specialized services. The only effect of doing so would be to confuse consumers, who would have the benefit of such policies when using their desktop (wired connection) or their laptop (WiFi connection) but not their iPhone or Android phone (frequently, but not always, CMRS connection) – and to leave substantial public interest benefits on the table.

Finally, there is no basis to conclude that the technical challenges associated with running a broadband Internet access network using CMRS frequencies and technologies are different in any way that matters with respect to the Commission's policies regarding "openness," including any rules or policies the Commission may establish regarding specialized services. To the contrary, the challenges Bright House Networks and other "wireline" providers face when confronting rapidly increasing consumer demand are directly parallel to those facing traditional "wireless" networks confronting the same situation. We can either find a way to get more

---

[65] We recognize that in some situations, some providers may use point-to-point microwave links, rather than carrier private line services, for this "backhaul" function. But by the same token, non-CMRS carriers routinely use point-to-point microwave services in some circumstances as well.

bandwidth over our existing fiber-coax network with its existing set of nodes, or we can reuse our current bandwidth more intensely by creating a new array of more numerous, but smaller nodes. In a directly parallel fashion, a traditional wireless provider faced with capacity constraints can either obtain more spectrum for use with its existing set of cell sites, or it can reuse its existing spectrum more intensely by creating a new array of more numerous, but smaller cells. So, putting aside the fact that there are few if any networks that are purely "wireline" or "wireless," from the perspective of managing a multipurpose network in the face of rising consumer demand, there is no meaningful distinction between the challenges faced by a prototypical "wired" network and a prototypical "wireless" CMRS network.[66]

For all these reasons, it would be a mistake to establish one set of rules for wired Internet access and a different set of rules – or no rules – for wireless access.[67] Whatever technical differences may characterize a traditional "wireless" network as opposed to other retail broadband Internet access networks, those differences are irrelevant to the policy concerns now before the Commission.

V. CONCLUSION

For the reasons stated above, the Commission should not impose any new rules or requirements with regard to specialized services offered by retail network providers. There is simply no evidence that there is a problem to be solved here. To the contrary, the evidence

---

[66] *See* Bright House Networks, LLC Comments in GN Docket No. 10-127 (July 15, 2010) at 14 n.31. Probably the only actually distinctive feature of a traditional "wireless" network is the need to be able to offer users seamless handoff from cell site to cell site while traveling at considerable speeds. This specific technical issue, however, has no bearing whatsoever on the Commission's policies regarding "openness."

[67] This applies "both ways" – while there is no reason to exempt CMRS-based networks from any Internet openness and nondiscrimination policies the Commission might apply to providers of broadband Internet access, there is no more reason to impose rules or restrictions regarding specialized services on CMRS-based networks than there is to impose such rules on cable- or telephone-company-based networks. In each case, the same policies and rules should apply regardless of network technology.

shows that it makes good technical sense to provide some services in a "specialized" way, and there is no reason to restrict the ability of retail network providers to work with rights holders to develop new business or technical models for delivering content and applications, whether or not it might be possible in some way to deliver them via the Internet. Moreover, because the Internet infrastructure is provided by many different entities in common, all members of that commons have well-established incentives to work cooperatively with each other to resolve any issues surrounding specialized services that might exist, making Commission action both unnecessary and potentially harmful in the highly technical area of Internet network and traffic engineering.

Finally, there is no sound technical or policy reason to treat wireless networks any differently than wired networks, whether with respect to any general "openness" or nondiscrimination obligations that might exist for broadband Internet access, or with respect to specialized services.

Respectfully submitted,

Marva Johnson Vice President Technology Policy and Industry Affairs Bright House Networks, LLC	Paul Glist Christopher W. Savage Davis Wright Tremaine LLP 1919 Pennsylvania Ave. NW Washington, D.C. 20006-3402 Arthur J. Steinhauer Cody Harrison Sabin, Bermant & Gould LLP Four Times Square New York, NY 10036 Counsel for Bright House Networks, LLC

October 12, 2010